The Rare Eclipse Problem on Tiles: Quantised Embeddings of Disjoint Convex Sets

Valerio Cambareri, Chunlei Xu, Laurent Jacques
ISPGroup, ICTEAM/ELEN, Université catholique de Louvain, Louvain-la-Neuve, Belgium.

arXiv:1702.04664v1 [cs.IT] 15 Feb 2017

***Abstract*—Quantised random embeddings are an efficient dimensionality reduction technique which preserves the distances of low-complexity signals up to some controllable additive and multiplicative distortions. In this work, we instead focus on verifying when this technique preserves the separability of two disjoint closed convex sets, *i.e.*, in a quantised view of the "rare eclipse problem" introduced by Bandeira *et al.* in 2014. This separability would ensure exact classification of signals in such sets from the signatures output by this non-linear dimensionality reduction. We here present a result relating the embedding's dimension, its quantiser resolution and the sets' separation, as well as some numerically testable conditions to illustrate it. Experimental evidence is then provided in the special case of two ℓ_2-balls, tracing the phase transition curves that ensure these sets' separability in the embedded domain.**

***Index Terms*—Random embeddings, dimensionality reduction, quantisation, compressive classification, phase transition.**

I. Introduction

Dimensionality reduction methods are a crucial part of very large-scale machine learning frameworks, as they are in charge of mapping (with negligible losses) the information contained in high-dimensional data to a low-dimensional domain, thus minimising the computational effort of learning tasks. We here focus on a class of *non-linear*, *non-adaptive* dimensionality reduction methods, *i.e.*, *quantised random embeddings*, as obtained[1] by applying to $\boldsymbol{x} \in \mathcal{K}$ (with $\mathcal{K} \subset \mathbb{R}^n$ any *dataset*)

$$\boldsymbol{y} = \mathsf{A}(\boldsymbol{x}) \coloneqq \mathcal{Q}_\delta(\boldsymbol{\Phi}\boldsymbol{x} + \boldsymbol{\xi}), \tag{1}$$

where $\boldsymbol{\Phi} \in \mathbb{R}^{m\times n}$ is a Gaussian random *sensing matrix*, *i.e.*, $\boldsymbol{\Phi} \sim \mathcal{N}^{m\times n}(0,1)$; $\mathcal{Q}_\delta(\cdot) \coloneqq \delta\lfloor \frac{\cdot}{\delta} \rfloor$ is a *uniform scalar quantiser* of *resolution* $\delta > 0$ (applied component-wise), yielding a *signature* $\boldsymbol{y} \in \delta\mathbb{Z}^m$; $\boldsymbol{\xi} \sim \mathcal{U}^m([0,\delta])$ is some *dither* drawn uniformly in $[0,\delta]^m$, which is fundamental to stabilise the action of the quantiser [1], [2].

The non-linear map described by (1) produces *compact* signatures $\boldsymbol{y}$, either in terms of dimension $m \ll n$, or of bits per entry (controlled by δ) even if $m > n$ [3]. Learning tasks such as classification may then run on $\boldsymbol{y} \in \mathsf{A}(\mathcal{K})$ rather than $\boldsymbol{x} \in \mathcal{K}$ at reduced storage, transmission, and computational costs, with accuracy depending on m, δ. However, contrarily to other non-linear maps (*e.g.*, [4], [5]), (1) retains *quasi-isometry* properties [2], [6] that grant, under some requirements on m (*i.e.*, *sample complexity* bounds), the *recovery* of $\boldsymbol{x}$ from $\boldsymbol{y}$ using appropriate algorithms [7].

E-mail: {valerio.cambareri, chunlei.xu, laurent.jacques}@uclouvain.be.
The authors are partly funded by the Belgian National Fund for Scientific Research (FNRS) under the M.I.S.-FNRS project AlterSense. All authors have equally contributed to the realisation of this paper.

[1]Our notation conventions are reported at the end of this section.

In this contribution we aim to prove that generic learning tasks can run seamlessly on $\mathsf{A}(\mathcal{K})$ by the ability of (1) to preserve the *separability* of different *classes* in $\mathcal{K}$. These classes are described by Q disjoint closed convex sets, *i.e.*, $\mathcal{C}_i \subset \mathcal{K}, i \in [Q]$ so that $\forall i,j \in [Q], i \neq j, \mathcal{C}_i \cap \mathcal{C}_j = \emptyset$. Hence, we inquire whether testing if our data $\boldsymbol{x} \in \mathcal{C}_i$ is *equivalent* to doing so given $\boldsymbol{y}$ in (1); for this to hold, it is necessary that the classes' images $\mathsf{A}(\mathcal{C}_i), i \in [Q]$ are still *separable*, *i.e.*, $\forall i,j \in [Q], i \neq j$, $\mathsf{A}(\mathcal{C}_i) \cap \mathsf{A}(\mathcal{C}_j) = \emptyset$. If this is violated then no learning algorithm can perform exact classification, as the images would "eclipse" each other. This perspective builds upon that of Bandeira *et al.* [8], who introduced this *rare eclipse problem* for linear embeddings, as reviewed in Sec. II-A. Focusing on $Q = 2$ classes, in Sec. II-B we define the *quantised eclipse problem* and present our main result, *i.e.*, a sample complexity bound which states the conditions on m, δ, $\boldsymbol{\Phi}$, and $\mathcal{C}_i, i \in [Q]$ under which the images $\mathsf{A}(\mathcal{C}_i)$ are separable with high probability (*w.h.p.*). In Sec. II-C this is simplified by lower bounds to the latter probability, which have the advantage of being numerically testable for disjoint convex sets by solving convex optimisation problems. Among such sets, we detail the specific case of two high-dimensional ℓ_2-balls in Sec. II-D; this is explored numerically in Sec. III by computing phase transition curves on the above probability bound, indicating a regime with respect to (*w.r.t.*) m, δ for (1) in which the sets' separability is preserved.

Notation: Given a random variable (*r.v.*) X (*e.g.*, normal $\mathcal{N}(0,1)$ or uniform $\mathcal{U}([0,\delta])$ *r.v.'s*), we write $\boldsymbol{U} \sim X^{d_1\times d_2}$ (*e.g.*, $\mathcal{N}^{d_1\times d_2}(0,1)$) to denote the $d_1 \times d_2$ matrix (or vector, if $d_2 = 1$) with independent and identically distributed (*i.i.d.*) entries $U_{ij} \sim_{\text{i.i.d.}} X$. Spheres and balls in $\ell_p(\mathbb{R}^q)$ are denoted by $\mathbb{S}_p^{q-1}$ and $\mathbb{B}_p^q$. For a set $\mathcal{C} \subset \mathbb{R}^n$, its Chebyshev radius is $\mathrm{rad}(\mathcal{C}) = \inf\{r > 0 : \exists \boldsymbol{c} \in \mathbb{R}^n,\ \mathcal{C} \subset \boldsymbol{c} + r\mathbb{B}^n\}$; its image under a map A is $\mathsf{A}(\mathcal{C})$; its projection by a matrix $\boldsymbol{B}$ is $\boldsymbol{B}\mathcal{C}$. The cardinality of a set $\mathcal{C}$ reads $|\mathcal{C}|$, and $[Q] \coloneqq \{1,\dots,Q\}$. We denote by C, c, c', c'' constants whose value can change between lines. We also write $f \lesssim g$ if $\exists c > 0$ such that $f \leqslant c\, g$, and correspondingly for $f \gtrsim g$. Moreover, $f \simeq g$ means that $f \lesssim g$ and $g \lesssim f$.

Relation to Prior Work: Many contributions have discussed linear dimensionality reduction by $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}$ with $\boldsymbol{\Phi} \sim X^{m\times n}$ a random matrix having *i.i.d.* entries distributed as a sub-Gaussian *r.v.* X (for a survey, see [9]), *i.e.*, *random projections*. Following the work of Johnson and Lindenstrauss [10], such linear embeddings were soon recognised [11], [12] as distance-

preserving, non-adaptive[2] dimensionality reductions for *finite datasets*, *i.e.*, with $|\mathcal{K}| < \infty$. Moreover, several non-linear random embeddings are now available for more general models of $\mathcal{K}$ [2], [5], [13]–[16]; most results on such embeddings rely on preserving distances, rather than the separation between classes within $\mathcal{K}$. Regarding this last aspect, Dasgupta [17] first analysed the separability of a mixture-of-Gaussians dataset $\mathcal{K}$ after random projections. Later, with the rise of Compressed Sensing (CS), random projections followed by classification tasks were dubbed *compressive classification*. Davenport *et al.* [18] showed that if $\boldsymbol{\Phi}$ verifies the Restricted Isometry Property (RIP) *w.r.t.* a dataset $\mathcal{K}$ (*i.e.*, a *stable embedding*) then exact classification can be achieved on $\boldsymbol{\Phi}\mathcal{K}$ thanks to distance preservation; $\mathcal{K}$ was therein taken as a finite set, or the set of sparse signals. Reboredo *et al.* [19], [20] studied the limits of compressive classification in a Bayesian framework. Finally, Bandeira *et al.* [8] first explored with the tools of high-dimensional geometry the conditions for the separability of closed convex sets $\mathcal{C}_1, \mathcal{C}_2 \subset \mathcal{K} \coloneqq \mathbb{R}^n$ after random projections. We here extend their approach to quantised random embeddings given by (1) which, due to their non-linearity, is a non-trivial endeavour that is currently lacking in the literature.

II. Quantised Random Embeddings and the Rare Eclipse Problem

A. The Rare Eclipse Problem

Let us first recall the fundamental question introduced by Bandeira *et al.* [8] and their main result as follows.

Problem 1 (Rare Eclipse Problem (from [8]))**.** *Let $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ be closed convex sets, $\boldsymbol{\Phi} \sim \mathcal{N}^{m\times n}(0,1)$. Given $\eta \in (0,1)$, find the smallest m so that*

$$p_0 \coloneqq \mathbb{P}[\boldsymbol{\Phi}\mathcal{C}_1 \cap \boldsymbol{\Phi}\mathcal{C}_2 = \emptyset] \geqslant 1-\eta. \tag{2}$$

Prob. 1 is equivalent to ensuring, for all $\boldsymbol{x}_1 \in \mathcal{C}_1$, $\boldsymbol{x}_2 \in \mathcal{C}_2$, that $\boldsymbol{\Phi}\boldsymbol{x}_1 \neq \boldsymbol{\Phi}\boldsymbol{x}_2$ with $\boldsymbol{\Phi} \sim \mathcal{N}^{m\times n}(0,1)$. Let us define the *difference set* $\mathcal{C}^- \coloneqq \mathcal{C}_1 - \mathcal{C}_2 = \{\boldsymbol{z} \coloneqq \boldsymbol{x}_1 - \boldsymbol{x}_2 : \boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2\}$. We can then cast (2) in terms of the kernel of $\boldsymbol{\Phi}$, *i.e.*,

$$\begin{aligned} p_0 &= \mathbb{P}[\forall \boldsymbol{z} \in \mathcal{C}^-, \boldsymbol{\Phi}\boldsymbol{z} \neq \boldsymbol{0}_m] = 1 - \mathbb{P}[\exists \boldsymbol{z} \in \mathcal{C}^- : \boldsymbol{\Phi}\boldsymbol{z} = \boldsymbol{0}_m] \\ &\geqslant 1-\eta, \textit{ i.e.}, \mathbb{P}[\mathrm{Ker}(\boldsymbol{\Phi}) \cap \mathcal{C}^- \neq \emptyset] \leqslant \eta. \end{aligned} \tag{3}$$

Intuitively, η in (2) will increase with the "size" of $\mathcal{C}^-$, as its intersection with $\mathrm{Ker}(\boldsymbol{\Phi})$ will be non-empty. This size is here measured by the *Gaussian mean width*, *i.e.*, for any set $\mathcal{C}$,

$$w(\mathcal{C}) \coloneqq \mathbb{E}_{\boldsymbol{g}} \sup_{\boldsymbol{x}\in\mathcal{C}} |\boldsymbol{g}^\top \boldsymbol{x}|, \; \boldsymbol{g} \sim \mathcal{N}^n(0,1).$$

Bandeira *et al.* then realised that (3) is found by Gordon's Escape Theorem [21] since, by arbitrarily scaling $\mathcal{C}^-$ that amounts to taking the cone $\mathbb{R}_+\mathcal{C}^-$, and by its intersection with the sphere $\mathbb{S}_2^{n-1}$, we obtain a *mesh* (*i.e.*, a closed subset of $\mathbb{S}_2^{n-1}$). Let us then define $\mathcal{S} \coloneqq (\mathbb{R}_+\mathcal{C}^-) \cap \mathbb{S}_2^{n-1}$ of width $w_\cap \coloneqq w(\mathcal{S})$, and report their main result (its proof is in [8]).

[2]Not requiring any potentially large or unavailable training dataset, as opposed to, *e.g.*, principal component analysis

Proposition 1 (Corollary 3.1 in [8])**.** *In the setup of Prob. 1, given $\eta \in (0,1)$, if $m \gtrsim (w_\cap + \sqrt{2\log\frac{1}{\eta}})^2 + 1$ then $p_0 \geqslant 1-\eta$.*

Hence, the sample complexity of Prob. 1 is sharply characterised for any difference set whose $w_\cap$ is given or bounded.

B. The Quantised Eclipse Problem

Extending Prop. 1 to quantised random embeddings as in (1) is not simple. To begin with, any two closed convex sets $\mathcal{C}_1, \mathcal{C}_2$ would now be mapped into two *countable* sets $\mathsf{A}(\mathcal{C}_1), \mathsf{A}(\mathcal{C}_2) \subset \delta\mathbb{Z}^m$; verifying when they "collide" is our key question below.

Problem 2 (Quantised Eclipse Problem)**.** *Let $\mathcal{C}_1, \mathcal{C}_2 \subset \mathbb{R}^n : \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ be closed convex sets, and A defined in (1) with $\delta > 0$. Given $\eta \in (0,1)$, find the smallest m so that*

$$p_\delta \coloneqq \mathbb{P}[\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) = \emptyset] \geqslant 1-\eta. \tag{4}$$

Note that, since A itself uses $\boldsymbol{\Phi} \sim \mathcal{N}^{m\times n}(0,1)$ before quantisation, $\boldsymbol{\Phi}\mathcal{C}_1 \cap \boldsymbol{\Phi}\mathcal{C}_2 \neq \emptyset \implies \mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) \neq \emptyset$; hence, $p_0 \geqslant p_\delta$ given the same $\boldsymbol{\Phi}, \mathcal{C}_1, \mathcal{C}_2$. However, the converse does not hold since $\Phi\mathcal{C}_1 \cap \Phi\mathcal{C}_2 = \emptyset$ by itself does not suffice to ensure $\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) = \emptyset$ due to, *e.g.*, coarse quantisation with large δ or some draws of $\boldsymbol{\xi}$ in (1). Then, letting the event $\mathsf{E} \coloneqq \{\forall \boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2, \mathsf{A}(\boldsymbol{x}_1) \neq \mathsf{A}(\boldsymbol{x}_2)\}$, we see (4) equals

$$p_\delta = \mathbb{P}[\mathsf{E}] \geqslant 1-\eta, \textit{ i.e.}, \mathbb{P}[\mathsf{E}^\mathrm{c}] \leqslant \eta, \tag{5}$$

$$\mathsf{E}^\mathrm{c} \coloneqq \{\exists\ \boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2 : \mathsf{A}(\boldsymbol{x}_1) = \mathsf{A}(\boldsymbol{x}_2)\} \tag{6}$$

Hence, η bounds the probability that any two $\boldsymbol{x}_1 \in \mathcal{C}_1$, $\boldsymbol{x}_2 \in \mathcal{C}_2$ are *consistent*. Note that, by consistency [2], $\mathsf{A}(\boldsymbol{x}_1) = \mathsf{A}(\boldsymbol{x}_2) \implies \|\boldsymbol{\Phi}\boldsymbol{z}\|_\infty < \delta$ with $\boldsymbol{z} = \boldsymbol{x}_1 - \boldsymbol{x}_2 \in \mathcal{C}^-$. Thus, introducing the *separation* $\sigma \coloneqq \min_{\boldsymbol{z}\in\mathcal{C}^-} \|\boldsymbol{z}\|_2$, it is expected that η will decay to 0 as σ increases and δ decreases.

This is also sustained by the fact that A is known to respect *w.h.p.* the Quantised Restricted Isometry Property (QRIP) [6] over some $\mathcal{K} \subset \mathbb{R}^n$ provided $\boldsymbol{\Phi}$ satisfies a (ℓ_1, ℓ_2)-form of the RIP (see Lemma 1) and m is large before the dimension of $\mathcal{K}$. If the QRIP holds, we would then have, for all $\boldsymbol{u}_1, \boldsymbol{u}_2 \in \mathcal{K}$,

$$|\tfrac{1}{m}\|\mathsf{A}(\boldsymbol{u}_1) - \mathsf{A}(\boldsymbol{u}_2)\|_1 - c'\|\boldsymbol{u}_1 - \boldsymbol{u}_2\|_2| \leqslant \varepsilon\|\boldsymbol{u}_1 - \boldsymbol{u}_2\|_2 + c\delta\varepsilon',$$

for some controllable distortions $\varepsilon, \varepsilon' > 0$ and constants $c, c' > 0$. With $\mathcal{K} \coloneqq \mathcal{C}_1 \cup \mathcal{C}_2$, $\boldsymbol{x}_1 \in \mathcal{C}_1$ and $\boldsymbol{x}_2 \in \mathcal{C}_2$, this ensures that $\frac{1}{m}\|\mathsf{A}(\boldsymbol{x}_1) - \mathsf{A}(\boldsymbol{x}_2)\|_1 \geqslant (c'-\varepsilon)\|\boldsymbol{z}\|_2 - c\delta\varepsilon' \geqslant (c'-\varepsilon)\sigma - c\delta\varepsilon'$. Thus, $\mathsf{A}(\boldsymbol{x}_1) \neq \mathsf{A}(\boldsymbol{x}_2)$ simply follows if $\frac{\sigma}{\delta} > \frac{c\varepsilon'}{c'-\varepsilon}$.

Before introducing our main result, let us present two lemmata, whose proof is given in the Appendix. The first assesses when $\boldsymbol{\Phi} \sim \mathcal{N}^{m\times n}(0,1)$ respects a (ℓ_1, ℓ_2)-form of the RIP for a mesh (see, *e.g.*, [15, Cor. 2.3], [22]).

Lemma 1. *Let $\epsilon_0 > 0$ and $\mathcal{S} \subset \mathbb{S}_2^{n-1}$. If $m \gtrsim \epsilon_0^{-2} w^2(\mathcal{S})$ and $\boldsymbol{\Phi} \sim \mathcal{N}^{m\times n}(0,1)$, then there exist some $C, c > 0$ such that, with probability exceeding $1 - C\exp(-c\epsilon_0^2 m)$ and $\kappa_0 = \sqrt{\frac{\pi}{2}}$,*

$$(1-\epsilon_0) \leqslant \tfrac{\kappa_0}{m}\|\boldsymbol{\Phi}\boldsymbol{u}\|_1 \leqslant (1+\epsilon_0), \quad \forall \boldsymbol{u} \in \mathcal{S}. \tag{7}$$

Thus, provided $m \gtrsim \epsilon_0^{-2} w_\cap^2$ and defining $\mathcal{D}_{\ell_1}(\boldsymbol{a}, \boldsymbol{b}) \coloneqq \frac{1}{m}\|\boldsymbol{a} - \boldsymbol{b}\|_1$, applying Lemma 1 to $\mathcal{S} \coloneqq (\mathbb{R}_+\mathcal{C}^-) \cap \mathbb{S}_2^{n-1}$ yields

$$\left|\kappa_0 \mathcal{D}_{\ell_1}(\boldsymbol{\Phi}\boldsymbol{x}_1, \boldsymbol{\Phi}\boldsymbol{x}_2) - \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2\right| \leqslant \epsilon_0 \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2, \tag{8}$$

with the same probability and for all $\boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2$, since $\frac{\boldsymbol{x}_1-\boldsymbol{x}_2}{\|\boldsymbol{x}_1-\boldsymbol{x}_2\|_2} \in \mathcal{S}$. Moreover, since $w^2(\mathbb{S}_2^{n-1}) \lesssim n$, provided $m \gtrsim \epsilon^{-2} n$ for some $\epsilon > 0$, we also have with probability exceeding $1 - C\exp(-c\epsilon^2 m)$,

$$(1-\epsilon)\|\boldsymbol{u}\|_2 \leqslant \tfrac{\kappa_0}{m}\|\boldsymbol{\Phi}\boldsymbol{u}\|_1 \leqslant (1+\epsilon)\|\boldsymbol{u}\|_2, \ \forall \boldsymbol{u} \in \mathbb{R}^n. \quad (9)$$

The second lemma proves that the mapping $\mathsf{A}'(\cdot) \coloneqq \mathcal{Q}(\cdot + \boldsymbol{\xi})$, with $\boldsymbol{\xi} \sim \mathcal{U}^m([0,\delta])$, embeds[3] *w.h.p.* $\mathbb{R}^m$ in $\delta\mathbb{Z}^m$ in the metric $\mathcal{D}_{\ell_1}$ and up to some controlled distortions. This lemma uses the Kolmogorov entropy $\mathcal{H}_q(\mathcal{E},\rho) \coloneqq \log \mathcal{N}_q(\mathcal{E},\rho)$ of a bounded subset $\mathcal{E} \subset \mathbb{R}^m$ in the ℓ_q-metric ($q \geqslant 1$) defined for $\rho > 0$, with $\mathcal{N}_p(\mathcal{E},\rho)$ the cardinality of its smallest ρ-covering in the same metric.

Lemma 2. *Let $\mathcal{E} \subset \mathbb{R}^m$ be a bounded set. Given $\epsilon, \delta > 0$, if*

$$m \gtrsim \epsilon^{-2}\mathcal{H}_1(\mathcal{E}, \tfrac{m\delta\epsilon^2}{1+\epsilon}),$$

then, for $\boldsymbol{\xi} \sim \mathcal{U}^m([0,\delta])$ and with probability exceeding $1 - C\exp(-cm\epsilon^2)$ for some $C, c > 0$, we have

$$\left|\mathcal{D}_{\ell_1}(\mathsf{A}'(\boldsymbol{a}),\mathsf{A}'(\boldsymbol{b})) - \mathcal{D}_{\ell_1}(\boldsymbol{a},\boldsymbol{b})\right| \lesssim \delta\epsilon, \ \forall \boldsymbol{a},\boldsymbol{b} \in \mathcal{E}. \quad (10)$$

We are finally able to state our main result, solving Prob. 2.

Proposition 2. *In the setup of Prob. 2, let $r_i \coloneqq \mathrm{rad}(\mathcal{C}_i)$, $i \in \{1,2\}$, $r \coloneqq r_1 + r_2$, and A defined in (1) with $\delta > 0$. Given $\eta \in (0,1)$, if*

$$m \gtrsim (w_\cap^2 + n\tfrac{\delta^2}{\sigma^2})(1 + \log(1 + \tfrac{rm}{\delta n}) + w_\cap^{-2}\log\tfrac{1}{\eta}), \quad (11)$$

then $p_\delta \geqslant 1 - \eta$.

Proof of Prop. 2. Let us first observe when (10) holds with $\boldsymbol{\Phi} \sim \mathcal{N}^{m\times n}(0,1)$, $\mathcal{E} \coloneqq \boldsymbol{\Phi}\mathcal{C}_\cup$, $\mathcal{C}_\cup \coloneqq \mathcal{C}_1 \cup \mathcal{C}_2 \subset \mathbb{R}^n$. This will be useful later to characterise when $\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) = \emptyset$. Let $\mathcal{R}_\cup$ be a ρ-covering in the ℓ_2-metric of $\mathcal{C}_\cup$ for some $\rho > 0$ to be specified below. If $m \gtrsim \epsilon^{-2}n$ for some $\epsilon > 0$, we have from (9) that, with probability exceeding $1 - C\exp(-c\epsilon^2 m)$, the event E_0 where $\boldsymbol{\Phi}\mathcal{R}_\cup$ is a ρ'-covering of $\boldsymbol{\Phi}\mathcal{C}_\cup$ holds with $\rho' = 2m(1+\epsilon)\rho$. This proves that, conditionally to E_0 and for $\mathcal{E} = \boldsymbol{\Phi}\mathcal{C}_\cup$, $\mathcal{H}_1(\mathcal{E},\rho') \leqslant \mathcal{H}_2(\mathcal{C}_\cup,\rho)$. However, $\mathcal{H}_2(\mathcal{C}_\cup,\rho) \leqslant \log 2 + \max(\mathcal{H}_2(\mathcal{C}_1,\rho),\mathcal{H}_2(\mathcal{C}_2,\rho)) \lesssim \max(\mathcal{H}_2(\mathcal{C}_1,\rho),\mathcal{H}_2(\mathcal{C}_2,\rho))$. Moreover, we have $\mathcal{H}_2(\mathcal{C}_i,\rho) \lesssim n\log(1+\frac{r_i}{\rho})$ [23], so that $\mathcal{H}_2(\mathcal{C}_\cup,\rho) \lesssim n\log(1+\frac{r}{\rho})$. Setting $\rho' \coloneqq \frac{m\delta\epsilon^2}{1+\epsilon}$ gives $\rho = \frac{\delta\epsilon^2}{2(1+\epsilon)^2}$ and finally

$$\mathcal{H}_1(\boldsymbol{\Phi}\mathcal{C}_\cup,\rho) \lesssim n\log(1 + \tfrac{2r(1+\epsilon)^2}{\delta\epsilon^2}).$$

Consequently, conditionally to E_0 which only depends on $\boldsymbol{\Phi}$, Lemma 2 provides that if $m \gtrsim \epsilon^{-2}n\log(1+\frac{2r(1+\epsilon)^2}{\delta\epsilon^2})$ then, with probability exceeding $1 - C\exp(-cm\epsilon^2)$, we get the occurrence of a new event, E_0', where (10) holds with $\boldsymbol{a} = \boldsymbol{\Phi}\boldsymbol{u}$ and $\boldsymbol{b} = \boldsymbol{\Phi}\boldsymbol{v}$ for all $\boldsymbol{u},\boldsymbol{v} \in \mathcal{C}_\cup$. Under the same conditions, since $\mathbb{P}[\mathsf{E}_0'] \geqslant \mathbb{P}[\mathsf{E}_0'|\mathsf{E}_0]\mathbb{P}[\mathsf{E}_0]$, E_0' occurs unconditionally with $\mathbb{P}[\mathsf{E}_0'] \geqslant 1 - C'\exp(-c'm\epsilon^2)$, for some $C', c' > 0$.

Second, if $m \gtrsim \epsilon_0^{-2}w_\cap^2$ for some $\epsilon_0 > 0$, Lemma 1 states that the event E_1, where (8) is respected for all $\boldsymbol{x}_1 \in \mathcal{C}_1$ and all $\boldsymbol{x}_2 \in \mathcal{C}_2$, holds with probability exceeding $1 - C\exp(-c\epsilon_0^2 m)$.

[3] That is, in the Gromov-Hausdorff sense [15].

Given $\eta \in (0,1)$ and $\epsilon = \epsilon(\epsilon_0) \coloneqq \frac{\sqrt{n}}{w_\cap}\epsilon_0$, *i.e.*, with $\epsilon \gtrsim \epsilon_0$ since $w_\cap \lesssim \sqrt{n}$, the union bound yields that E_0 and E_1 jointly hold with probability exceeding $1 - C\exp(-c\epsilon_0^2 m) \geqslant 1 - \eta$ provided

$$m \gtrsim \epsilon_0^{-2}\big(w_\cap^2\,(1 + \log(1 + \tfrac{2r\,(1+\epsilon(\epsilon_0))^2}{\delta\,\epsilon(\epsilon_0)^2})) + \log\tfrac{1}{\eta}\big). \quad (12)$$

In this case, for all $\boldsymbol{u} \in \mathcal{C}_1$ and $\boldsymbol{v} \in \mathcal{C}_2$ (or *vice versa*), (10) (with $\boldsymbol{a} \coloneqq \boldsymbol{\Phi}\boldsymbol{u}$ and $\boldsymbol{b} \coloneqq \boldsymbol{\Phi}\boldsymbol{v}$) and (8) give, for some $c > 0$,

$$\begin{aligned}
&\mathcal{D}_{\ell_1}(\mathsf{A}(\boldsymbol{u}),\mathsf{A}(\boldsymbol{v})) = \mathcal{D}_{\ell_1}(\mathsf{A}'(\boldsymbol{\Phi}\boldsymbol{u}),\mathsf{A}'(\boldsymbol{\Phi}\boldsymbol{v}))\\
&\geqslant \mathcal{D}_{\ell_1}(\boldsymbol{\Phi}\boldsymbol{u},\boldsymbol{\Phi}\boldsymbol{v}) - c\delta\epsilon \geqslant \kappa_0^{-1}\,(1-\epsilon_0)\|\boldsymbol{u}-\boldsymbol{v}\|_2 - c\delta\epsilon\\
&\geqslant \kappa_0^{-1}\,(1-\epsilon_0)\sigma - c\delta\epsilon = \kappa_0^{-1}\,(1-\epsilon_0)\sigma - \tfrac{c\delta\sqrt{n}}{w_\cap}\epsilon_0.
\end{aligned}$$

In order to have $\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) = \emptyset$, the last quantity must be positive. Since $\epsilon_0 > 0$, this clearly happens if $\kappa_0^{-1}\,(1-\epsilon_0)\sigma - \frac{c\delta\sqrt{n}}{w_\cap}\epsilon_0 = \frac{c\delta\sqrt{n}}{w_\cap}\epsilon_0$, which gives

$$\epsilon_0^{-2}w_\cap^2 = (w_\cap + 2c\kappa_0\sqrt{n}\tfrac{\delta}{\sigma})^2 \lesssim w_\cap^2 + n\tfrac{\delta^2}{\sigma^2}.$$

Moreover, from the value of $\epsilon = \epsilon(\epsilon_0)$ set above, $\frac{2r(1+\epsilon)^2}{\delta\epsilon^2} \leqslant 2\frac{r}{n\delta}(w_\cap^2 + n\frac{\delta^2}{\sigma^2})$, so that (12) is satisfied if (11) holds. This gives finally that $p_\delta \geqslant 1 - \eta$ under this condition. □

Interestingly, up to diverging $\log$ factors (possibly due to proof artefacts), the requirement of Prop. 1 can be seen as a special case of (11) when $\delta \to 0^+$, *i.e.*, for a "vanishing" quantiser, since $(w_\cap + \sqrt{2\log\frac{1}{\eta}})^2 \lesssim w_\cap^2 + \log\frac{1}{\eta}$. Finally, the application of Prop. 2 to more than two sets is possible, and will be included in an extended version of this paper.

C. Testable Conditions by Convex Problems

To properly verify the bound on p_δ in Prop. 2 we should test the existence of any element in $\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) \subset \delta\mathbb{Z}^m$, *i.e.*, of any two *consistent* vectors $\boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2 : \mathsf{A}(\boldsymbol{x}_1) = \mathsf{A}(\boldsymbol{x}_2)$. As expected, this search is computationally intractable, so we now deduce numerically testable, albeit less tight lower bounds for p_δ. Let us first define the *consistency margin* τ

$$\tau \coloneqq \min_{\boldsymbol{z}\in\mathcal{C}^-} \|\boldsymbol{\Phi}\boldsymbol{z}\|_\infty, \quad (13)$$

that is a function of $\boldsymbol{\Phi}$ and $\mathcal{C}^-$, and can be related to the minimal separation σ as defined above. Moreover, the event $\tau > \delta$ depends only on $\boldsymbol{\Phi}$, so we can write (5) as

$$p_\delta = \mathbb{P}_{\boldsymbol{\Phi},\boldsymbol{\xi}}[\mathsf{E}|\tau \leqslant \delta]\mathbb{P}_{\boldsymbol{\Phi}}[\tau \leqslant \delta] + \bar{p}_\delta \geqslant \bar{p}_\delta, \quad (14)$$

where $\bar{p}_\delta \coloneqq \mathbb{P}_{\boldsymbol{\Phi}}[\tau > \delta]$ (*i.e.*, $\mathbb{P}_{\boldsymbol{\Phi},\boldsymbol{\xi}}[\mathsf{E}|\tau > \delta] = 1$) since $\|\boldsymbol{\Phi}(\boldsymbol{x}_1 - \boldsymbol{x}_2)\|_\infty \geqslant \tau > \delta \implies \mathsf{A}(\boldsymbol{x}_1) \neq \mathsf{A}(\boldsymbol{x}_2)$, while the converse does not hold. Note that $\bar{p}_\delta$ fully accounts for the cases in which $\mathrm{Ker}(\boldsymbol{\Phi}) \cap \mathcal{C}^- \neq \emptyset$, since if $\boldsymbol{\Phi}\boldsymbol{z} = \boldsymbol{0}_m \implies \tau = 0$. Clearly, we can now estimate $\bar{p}_\delta$ as τ can be computed for each $\boldsymbol{\Phi}$ when the optimisation problem (13) is convex (*i.e.*, iff $\mathcal{C}^-$ is, as for disjoint convex sets).

To tighten this bound and fully leverage dithering, we form a partition $\cup_j \mathcal{C}^{(j)} = \mathcal{C}^-$ formed by the cones

$$\mathcal{C}^{(j)} \coloneqq \{\boldsymbol{z} \in \mathcal{C}^- : |\boldsymbol{\varphi}_j^\top \boldsymbol{z}| \geqslant |\boldsymbol{\varphi}_i^\top \boldsymbol{z}|, \forall i \neq j \in [m]\} \subset \mathcal{C}^-.$$

We can now define for $j \in [m], \tau_j \coloneqq \min_{\boldsymbol{z}\in\mathcal{C}^{(j)}} |\boldsymbol{\varphi}_j^\top \boldsymbol{z}|$, where clearly $\tau_j \geqslant \tau$. Letting $\mathsf{A}_i(\boldsymbol{x}) \coloneqq \mathcal{Q}_\delta(\boldsymbol{\varphi}_i^\top \boldsymbol{z} + \xi_i)$, we use a

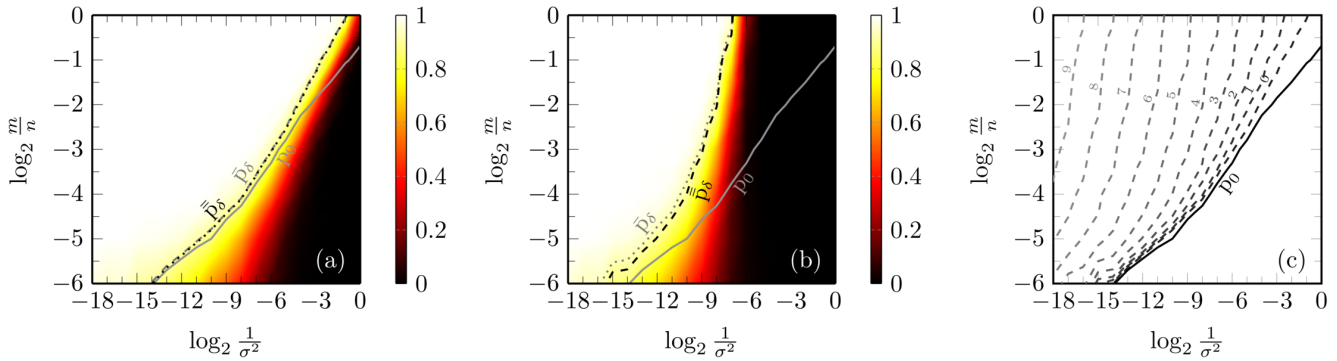

Figure 1. Empirical phase transitions of the quantised eclipse problem for the case of two disjoint ℓ_2-balls; for 128 random instances of $\mathbf{\Phi}$ and as a function of $\frac{1}{\sigma^2}$ and the rate $\frac{m}{n}$, we report for $\delta = 2^1$ (a) and $\delta = 2^4$ (b) the empirical estimate of $\bar{\bar{p}}_\delta$ (heat map), along with the phase transition curves at probability 0.9; (c) illustrates how the phase transition curves $\bar{\bar{\mathrm{p}}}_\delta$ vary for several values of δ, with $\log_2 \delta$ being annotated on each line.

shorthand for the event $\mathsf{E}_i \coloneqq \{\mathsf{A}_i(\boldsymbol{x}_1) \neq \mathsf{A}_i(\boldsymbol{x}_2)\}$ and bound

$$
\begin{aligned}
p_\delta &= \mathbb{P}_{\mathbf{\Phi},\boldsymbol{\xi}}[\forall \boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2, \exists i \in [m] : \mathsf{E}_i] \\
&= \mathbb{P}_{\mathbf{\Phi},\boldsymbol{\xi}}[\forall j \in [m], \boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2, \boldsymbol{x}_1 - \boldsymbol{x}_2 \in \mathcal{C}^{(j)}, \exists i \in [m] : \mathsf{E}_i] \\
&\geqslant \mathbb{P}_{\mathbf{\Phi},\boldsymbol{\xi}}[\forall j \in [m], \boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2, \boldsymbol{x}_1 - \boldsymbol{x}_2 \in \mathcal{C}^{(j)}, \mathsf{E}_j] \\
&= \mathbb{E}_{\mathbf{\Phi}} \mathbb{P}_{\boldsymbol{\xi}}[\forall j \in [m], \boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2, \boldsymbol{x}_1 - \boldsymbol{x}_2 \in \mathcal{C}^{(j)}, \mathsf{E}_j | \mathbf{\Phi}]
\end{aligned}
$$

Then, since the entries ξ_j of $\boldsymbol{\xi} \sim \mathcal{U}^m([0,\delta])$ are *i.i.d.*,

$$
\begin{aligned}
&\mathbb{P}_{\boldsymbol{\xi}}[\forall j \in [m], \boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2, \boldsymbol{x}_1 - \boldsymbol{x}_2 \in \mathcal{C}^{(j)}, \mathsf{E}_j | \mathbf{\Phi}] \\
&= \textstyle\prod_{j \in [m]} \mathbb{P}_{\xi_j}[\forall \boldsymbol{x}_1 \in \mathcal{C}_1, \boldsymbol{x}_2 \in \mathcal{C}_2 : \boldsymbol{x}_1 - \boldsymbol{x}_2 \in \mathcal{C}^{(j)}, \mathsf{E}_j | \mathbf{\Phi}] \\
&= \textstyle\prod_{j \in [m]} \min\{1, \frac{\tau_j}{\delta}\} \\
&\implies p_\delta \geqslant \mathbb{E}_{\mathbf{\Phi}} \textstyle\prod_{j \in [m]} \min\{1, \frac{\tau_j}{\delta}\} =: \bar{\bar{p}}_\delta, \qquad (15)
\end{aligned}
$$

where the second last line follows since E_j occurs whenever, given two intervals $\boldsymbol{\varphi}_j^\top \mathcal{C}_1, \boldsymbol{\varphi}_j^\top \mathcal{C}_2 \subset \mathbb{R}$ that are τ_j far apart, a quantiser threshold $\delta\mathbb{Z} + \xi_j$ falls between them. Hence, this event is identical to having $\mathbb{P}[\xi_j \in [0, \tau_j]] = \min\{1, \tau_j/\delta\}$ since $\xi_j \sim \mathcal{U}([0,\delta])$. The computational complexity of estimating $\bar{\bar{p}}_\delta$ is similar to that of (14), while (15) is sharper, as it can be shown that $\bar{\bar{p}}_\delta \geqslant \bar{p}_\delta$. However, we expect both bounds to be somewhat loose *w.r.t.* the one in Prop. 2.

D. The Case of Two Disjoint ℓ_2-Balls

We now briefly focus on the case of two ℓ_2-balls $\mathcal{C}_1 \coloneqq r_1 \mathbb{B}_2^n + \boldsymbol{c}_1$ and $\mathcal{C}_2 \coloneqq r_2 \mathbb{B}_2^n + \boldsymbol{c}_2$, for which $\mathcal{C}^- = r\mathbb{B}_2^n + \boldsymbol{c}$ with $\boldsymbol{c} \coloneqq \boldsymbol{c}_1 - \boldsymbol{c}_2$ and $r \coloneqq r_1 + r_2$. It is then shown in [24, Prop. 4.3] that $w_\cap \lesssim \frac{r}{\|\boldsymbol{c}\|_2}\sqrt{n} \simeq \frac{\sqrt{n}}{\sigma}$ when $\sigma \gg r$ since $\sigma = \|\boldsymbol{c}\|_2 - r$. We can now compare the sample complexities in Prop. 1 and Prop. 2: up to some log and additive factors, we see that Prob. 2 has rate $\frac{m}{n} \gtrsim \frac{1}{\sigma^2}(1+\delta^2)$, while Prob. 1 only requires $\frac{m}{n} \gtrsim \frac{1}{\sigma^2}$, hence showing the effect of δ that we will illustrate in our numerical experiments below.

III. Numerical Experiments

We now test the special case of Sec. II-D by generating random instances of $\mathbf{\Phi} \sim \mathcal{N}^{m \times n}(0,1)$ and[4] $\mathcal{C}^-$, and computing the quantities $\tau_j, j \in [m]$ and τ for each instance, as specified in Sec. II-C. This allows us to empirically estimate $\bar{p}_\delta, \bar{\bar{p}}_\delta$ respectively in (14), (15) on 128 trials for each of the configurations $n = 2^6$ and $m \in [2^0, 2^6]$, and varying $\mathcal{C}^-$ by fixing $r = 2$ and taking $\sigma = \|\boldsymbol{c}\|_2 - r \in [2^0, 2^9]$. The estimated values of $\bar{\bar{p}}_\delta$ are then reported as heat maps in Fig. 1a,b along with the phase transition curves $\bar{\bar{\mathrm{p}}}_\delta \coloneqq \{\bar{\bar{p}}_\delta \geqslant 0.9\}$, $\bar{\mathrm{p}}_\delta \coloneqq \{\bar{p}_\delta \geqslant 0.9\}$, and the linear case of Prop. 1 $\mathrm{p}_0 \coloneqq \{p_0 \geqslant 0.9\}$, with p_0 being estimated as in [8]. Given $\tau_j, j \in [m]$ for all instances, we compute in Fig. 1c the phase transition curves corresponding to $\bar{\bar{p}}_\delta$ for several $\delta = \{2^0, 2^1, \dots, 2^9\}$. For each curve, the event $\mathsf{A}(\mathcal{C}_1) \cap \mathsf{A}(\mathcal{C}_2) = \emptyset$ holds with probability at least 0.9. These curves are indeed compatible with the fact that $\log_2 \frac{m}{n} \gtrsim \log_2 \frac{1}{\sigma^2} + \log_2(1+\delta^2)$ (up to log factors, and as concluded in Sec. II-D). However, we suspect that $\bar{\bar{p}}_\delta$ is still not sufficiently tight to approach our theoretical, albeit computationally intractable, bound on $p\delta$, and leave this improvement to a future investigation.

IV. Conclusion

The fundamental limits of learning tasks with embeddings are being tackled in several studies; our result illustrates the requirements for exact classification after quantised random embedding of two disjoint closed convex sets. As we only developed cases in which the datasets $\mathcal{K}$ are not specified as low complexity sets, we will discuss them in future works, *e.g.*, for the case of Q disjoint "clusters" of sparse signals $\mathcal{C}_i, i \in [Q]$.

V. Appendix

Proof of Lemma 2. We adapt the proof of [6, Prop. 1]. Given $\rho > 0$ to be fixed later, let $\mathcal{E}_\rho$ be a ρ-covering of $\mathcal{E}$ in the ℓ_1-metric, *i.e.*, for all $\boldsymbol{a} \in \mathcal{E}$ there exists $\boldsymbol{a}_0 \in \mathcal{E}_\rho$ such that $\|\boldsymbol{a} - \boldsymbol{a}_0\|_1 \leqslant \rho$. Notice that since $X(\boldsymbol{a}, \boldsymbol{b}) \coloneqq \mathcal{D}_{\ell_1}(\mathsf{A}'(\boldsymbol{a}), \mathsf{A}'(\boldsymbol{b})) = \frac{1}{m}\sum_i |\mathcal{Q}(a_i + \xi_i) - \mathcal{Q}(b_i + \xi_i)| = \frac{1}{m}\sum_i X_i$, with the *i.i.d.* sub-Gaussian *r.v.'s* X_i such that $\mathbb{E}X_i = |a_i - b_i|$ [2, App. A], one can easily prove the concentration of $X(\boldsymbol{a}, \boldsymbol{b})$ around $\mathbb{E}X(\boldsymbol{a}, \boldsymbol{b}) = \mathcal{D}_{\ell_1}(\boldsymbol{a}, \boldsymbol{b})$ both on a fixed pair $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{E}$ and, by union bound, for all $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{E}_\rho$ since there are no more than $(e^{\mathcal{H}_1(\mathcal{E},\rho)})^2$ such pairs in $\mathcal{E}_\rho$. Unfortunately, the discontinuity of the mapping A' prevents us to directly extend this over the full set $\mathcal{E}$ by a continuity argument applied to each neighbourhood of the covering.

[4] By uniformity of $\mathrm{Ker}(\mathbf{\Phi})$, $\mathbf{\Phi} \sim \mathcal{N}^{m \times n}(0,1)$ over the Grassmannian at the origin, it is legitimate to fix a randomly drawn direction $\frac{\boldsymbol{c}}{\|\boldsymbol{c}\|_2}$ for the simulations.

However, this situation can be overcome by *softening* the pseudo-distance $d(\cdot,\cdot) \coloneqq |\mathcal{Q}(\cdot) - \mathcal{Q}(\cdot)|$ composing X [2], [15]. We first note that $d(a,b) \coloneqq \delta \sum_{k\in\mathbb{Z}} \mathbb{I}_{\mathcal{S}}(a-k\delta, b-k\delta)$, where $\mathcal{S} = \{(a,b)\in\mathbb{R}^2 : ab<0\}$ and $\mathbb{I}_{\mathcal{C}}(a,b)$ is the indicator of $\mathcal{C}$ evaluated in (a,b), *i.e.*, it is equal to 1 if $(a,b)\in\mathcal{C}$ and 0 otherwise. In fact, $d(a,b) = \delta|(\delta\mathbb{Z})\cap[a,b]|$, with $|\cdot|$ the cardinality operator, showing that d/δ counts the number of *thresholds* in $\delta\mathbb{Z}$ that can be inserted between a and b.

Introducing the set $\mathcal{S}^t = \{(a,b)\in\mathbb{R}^2 : a < -t, b > t\} \cup \{(a,b)\in\mathbb{R}^2 : a > t, b < -t\}$ for $t\in\mathbb{R}$, with $\mathcal{S}^0 = \mathcal{S}$, we can define a soft version of d by

$$d^t(a,b) \coloneqq \delta \sum_{k\in\mathbb{Z}} \mathbb{I}_{\mathcal{S}^t}(a-k\delta, b-k\delta). \tag{16}$$

Thanks to $\mathcal{S}^t$, the value of t determines a set of forbidden (or relaxed) intervals $\delta\mathbb{Z}+[-|t|,|t|] = \{\,[k\delta-|t|, k\delta+|t|] : k\in\mathbb{Z}\}$ if $t>0$ (respectively $t<0$) of size $2|t|$ and centred on the quantiser thresholds in $\delta\mathbb{Z}$. For $t>0$ a threshold of $\delta\mathbb{Z}$ is not counted in $d^t(a,b)$ if a or b fall in its forbidden interval, whereas for $t<0$ a threshold that is not between a and b can be counted if a or b fall inside its relaxed interval.

By extension, we can also define $\mathcal{D}^t(\boldsymbol{a},\boldsymbol{b}) \coloneqq \frac{1}{m}\sum_i d^t(a_i,b_i)$ for $\boldsymbol{a},\boldsymbol{b}\in\mathbb{R}^m$, so that $\mathcal{D}^0(\boldsymbol{a},\boldsymbol{b}) = \mathcal{D}_{\ell_1}\big(\mathcal{Q}(\boldsymbol{a}),\mathcal{Q}(\boldsymbol{b})\big)$. Interestingly, this distance displays the following continuity property [2, Lemma 2]. For $\boldsymbol{a},\boldsymbol{b}\in\mathcal{E}$, and $\boldsymbol{a}_0,\boldsymbol{b}_0$ their respective closest points in $\mathcal{E}_\rho$ we have, for every $t\in\mathbb{R}$ and[5] $P>0$,

$$\mathcal{D}^t(\boldsymbol{a},\boldsymbol{b}) \geqslant \mathcal{D}^{t+\frac{\rho P}{m}}(\boldsymbol{a}_0,\boldsymbol{b}_0) - 8(\tfrac{\delta}{P}+\tfrac{\rho}{m}), \tag{17}$$

$$\mathcal{D}^t(\boldsymbol{a},\boldsymbol{b}) \leqslant \mathcal{D}^{t-\frac{\rho P}{m}}(\boldsymbol{a}_0,\boldsymbol{b}_0) + 8(\tfrac{\delta}{P}+\tfrac{\rho}{m}). \tag{18}$$

Moreover, for $\boldsymbol{\xi}\sim\mathcal{U}^m([0,\delta])$ and $\boldsymbol{a},\boldsymbol{b}$ fixed, $\mathcal{D}^t(\boldsymbol{a}+\boldsymbol{\xi},\boldsymbol{b}+\boldsymbol{\xi})$ concentrates around its mean which is close to $\mathcal{D}_{\ell_1}(\boldsymbol{a},\boldsymbol{b})$ [2, Lemma 3]. In fact, $|\mathbb{E}\mathcal{D}^t(\boldsymbol{a}+\boldsymbol{\xi},\boldsymbol{b}+\boldsymbol{\xi}) - \mathcal{D}_{\ell_1}(\boldsymbol{a},\boldsymbol{b})| \lesssim |t|$, so that for some $c>0$,

$$\mathbb{P}\big[|\mathcal{D}^t(\boldsymbol{a}+\boldsymbol{\xi},\boldsymbol{b}+\boldsymbol{\xi}) - \mathcal{D}_{\ell_1}(\boldsymbol{a},\boldsymbol{b})| > 4|t| + \epsilon(\delta+|t|)\big] \lesssim e^{-c\epsilon^2 m}.$$

Therefore, by union bound and for some $P>0$ to be fixed soon, if $m \gtrsim \epsilon^{-2}\mathcal{H}_1(\mathcal{E},\rho)$ then

$$\begin{aligned}|\mathcal{D}^{t\pm\frac{\rho P}{m}}(\boldsymbol{a}_0'+\boldsymbol{\xi},\boldsymbol{b}_0'+\boldsymbol{\xi}) - \mathcal{D}_{\ell_1}(\boldsymbol{a}_0',\boldsymbol{b}_0')| \\ \leqslant 4|t| + 4\tfrac{\rho P}{m} + \epsilon(\delta+|t|+\tfrac{\rho P}{m}),\ \forall \boldsymbol{a}_0',\boldsymbol{b}_0'\in\mathcal{E}_\rho,\end{aligned} \tag{19}$$

with probability exceeding $1-Ce^{-c\epsilon^2 m}$ for some $C,c>0$.

Consequently, for any $\boldsymbol{a},\boldsymbol{b}\in\mathcal{E}$ and $\boldsymbol{a}_0,\boldsymbol{b}_0$ their respective closest point in $\mathcal{E}_\rho$, using (18) combined with (19), and since the triangular inequality provides $\mathcal{D}_{\ell_1}(\boldsymbol{a}_0,\boldsymbol{b}_0) \leqslant \mathcal{D}_{\ell_1}(\boldsymbol{a},\boldsymbol{b}) + \frac{2\rho}{m}$, we have with the same probability and for some $c>0$,

$$\begin{aligned}&\mathcal{D}^t(\boldsymbol{a}+\boldsymbol{\xi},\boldsymbol{b}+\boldsymbol{\xi}) \leqslant \mathcal{D}^{t-\frac{\rho P}{m}}(\boldsymbol{a}_0+\boldsymbol{\xi},\boldsymbol{b}_0+\boldsymbol{\xi}) + 8(\tfrac{\delta}{P}+\tfrac{\rho}{m}) \\ &\leqslant \mathcal{D}_{\ell_1}(\boldsymbol{a}_0,\boldsymbol{b}_0) + 4|t| + \tfrac{4\rho P}{m} + \epsilon(\delta+|t|+\tfrac{\rho P}{m}) + 8(\tfrac{\delta}{P}+\tfrac{\rho}{m}) \\ &\leqslant \mathcal{D}_{\ell_1}(\boldsymbol{a},\boldsymbol{b}) + c\big(|t| + \tfrac{\rho(1+P(1+\epsilon))}{m} + \epsilon(\delta+|t|) + (\tfrac{\delta}{P}+\tfrac{\rho}{m})\big) \\ &\leqslant \mathcal{D}_{\ell_1}(\boldsymbol{a},\boldsymbol{b}) + c(|t|+\delta\epsilon),\end{aligned}$$

where we finally set the free parameters as $P^{-1}=\epsilon$ and $\rho = m\delta\frac{\epsilon^2}{1+\epsilon} < m\delta\min(\epsilon,\epsilon^2)$, giving $\rho P \leqslant m\delta\epsilon$ and $\frac{\rho}{m}\leqslant\delta\epsilon$. The lower bound is obtained similarly using (17) with the minus case of (19), and Prop. 2 is finally obtained with $t=0$. □

[5] In [2, Lemma 2], it is assumed $P\geqslant 1$ but nothing prevents $P>0$.

REFERENCES

[1] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2325–2383, 1998.

[2] L. Jacques, "Small width, low distortions: quasi-isometric embeddings with quantized sub-Gaussian random projections," *arXiv preprint arXiv:1504.06170*, 2015.

[3] P. T. Boufounos, L. Jacques, F. Krahmer, and R. Saab, "Quantization and compressive sensing," in *Compressed Sensing and its Applications*. Springer, 2015, pp. 193–237.

[4] A. Rahimi and B. Recht, "Random Features for Large-Scale Kernel Machines," in *Advances in Neural Information Processing Systems 20*, J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, Eds. Curran Associates, Inc., 2008, pp. 1177–1184.

[5] P. T. Boufounos, S. Rane, and H. Mansour, "Representation and Coding of Signal Geometry," *arXiv preprint arXiv:1512.07636*, 2015.

[6] L. Jacques and V. Cambareri, "Time for dithering: fast and quantized random embeddings via the restricted isometry property," *arXiv preprint arXiv:1607.00816*, 2016.

[7] A. Moshtaghpour, L. Jacques, V. Cambareri, K. Degraux, and C. De Vleeschouwer, "Consistent Basis Pursuit for Signal and Matrix Estimates in Quantized Compressed Sensing," *IEEE Signal Processing Letters*, vol. 23, no. 1, pp. 25–29, 2016.

[8] A. S. Bandeira, D. G. Mixon, and B. Recht, "Compressive classification and the rare eclipse problem," *arXiv preprint arXiv:1404.3203*, 2014.

[9] R. Vershynin, "Introduction to the non-asymptotic analysis of random matrices," in *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012, pp. 210–268.

[10] W. B. Johnson and J. Lindenstrauss, "Extensions of Lipschitz mappings into a Hilbert space," *Contemporary Mathematics*, vol. 26, no. 189-206, pp. 1–1, 1984.

[11] S. Dasgupta and A. Gupta, "An elementary proof of a theorem of Johnson and Lindenstrauss," *Random Structures & Algorithms*, vol. 22, no. 1, pp. 60–65, Jan. 2003.

[12] D. Achlioptas, "Database-friendly random projections: Johnson-Lindenstrauss with binary coins," *Journal of Computer and System Sciences*, vol. 66, no. 4, pp. 671–687, Jun. 2003.

[13] S. Oymak and B. Recht, "Near-Optimal Bounds for Binary Embeddings of Arbitrary Sets," *arXiv preprint arXiv:1512.04433*, 2015.

[14] L. Jacques, "A Quantized Johnson–Lindenstrauss lemma: The finding of Buffon's needle," *IEEE Transactions on Information Theory*, vol. 61, no. 9, pp. 5012–5027, 2015.

[15] Y. Plan and R. Vershynin, "Dimension reduction by random hyperplane tessellations," *Discrete & Computational Geometry*, vol. 51, no. 2, pp. 438–461, 2014.

[16] ——, "Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach," *IEEE Transactions on Information Theory*, vol. 59, no. 1, pp. 482–494, 2013.

[17] S. Dasgupta, "Learning mixtures of Gaussians," in *40th Annual Symposium on Foundations of Computer Science*, 1999, pp. 634–644.

[18] M. A. Davenport, P. T. Boufounos, M. B. Wakin, and R. G. Baraniuk, "Signal processing with compressive measurements," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 445–460, 2010.

[19] H. Reboredo, F. Renna, R. Calderbank, and M. Rodrigues, "Compressive classification," in *2013 IEEE International Symposium on Information Theory Proceedings (ISIT)*, Jul. 2013, pp. 674–678.

[20] H. Reboredo, F. Renna, R. Calderbank, and M. R. Rodrigues, "Bounds on the Number of Measurements for Reliable Compressive Classification," *IEEE Transactions on Signal Processing*, vol. 64, no. 22, pp. 5778–5793, 2016.

[21] Y. Gordon, "On Milman's inequality and random subspaces which escape through a mesh in $\mathbb{R}$ n," in *Geometric Aspects of Functional Analysis*. Springer, 1988, pp. 84–106.

[22] G. Schechtman, "Two observations regarding embedding subsets of Euclidean spaces in normed spaces," *Advances in Mathematics*, vol. 200, no. 1, pp. 125–135, 2006.

[23] G. Pisier, *The volume of convex bodies and Banach space geometry*. Cambridge University Press, 1999, vol. 94.

[24] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp, "Living on the edge: phase transitions in convex programs with random data," *Information and Inference: A Journal of the IMA*, vol. 3, no. 3, p. 224, 2014.